Improved Bounds on Minimax Regret under Logarithmic Loss via Self-Concordance

Blair Bilodeau [1,2,3] Dylan J. Foster [4] Daniel M. Roy [1,2,3]

[1]University of Toronto [2]Vector Institute [3]Institute for Advanced Study [4]Massachusetts Institute of Technology

Contribution Summary

- **Tighter upper bounds** on minimax regret under logarithmic loss for complex expert classes.
- First **truncation-free argument** which improves on previous best results.
- Easily optimized form of upper bound which **does not require chaining**.
- Characterize a **lower bound** using techniques from regret for square loss.

Online Learning and Minimax Regret

Traditional statistical learning analyzes data in a *batch* to produce a prediction function, which is used on future observations assumed to be generated i.i.d. from the training distribution. **Online learning is a framework for predicting future observations without any assumptions about the data generating process.**

For rounds $t = 1, \dots, n$:

- Environment supplies *context* $x_t \in \mathcal{X}$, which depends on the history;
- Player *predicts* $\hat{p}_t \in [0,1]$, a distribution on binary observations;
- Adversary generates *observation* $y_t \in \{0,1\}$;
- Player incurs *loss* $\ell_{\log}(\hat{p}_t, y_t) = -y_t \log(\hat{p}_t) - (1-y_t)\log(1-\hat{p}_t)$.

Observe that the loss corresponds to the *negative log-likelihood* of the observation under the predicted distribution.

In general, the player's cumulative loss grows super-linearly in n.

Performance is measured with respect to a class of *experts* $\mathcal{F} \subseteq [0,1]^{\mathcal{X}}$. The player's goal is to **compete against the best expert in hindsight**, which characterizes their *regret*:

$$R_n^{\log}(\hat{\mathbf{p}}; \mathcal{F}, \mathbf{x}, \mathbf{y}) = \sum_{t=1}^{n} \ell_{\log}(\hat{p}_t, y_t) - \inf_{f \in \mathcal{F}} \sum_{t=1}^{n} \ell_{\log}(f(x_t), y_t).$$

The *minimax regret* is an **algorithm-free concept** that measures how difficult an expert class is to learn over worst-case observations.

$$R_n^{\log}(\mathcal{F}) = \left\langle\!\!\left\langle \sup_{x_t} \inf_{\hat{p}_t} \sup_{y_t} \right\rangle\!\!\right\rangle_{t=1}^{n} R_n^{\log}(\hat{\mathbf{p}}; \mathcal{F}, \mathbf{x}, \mathbf{y}).$$

Goal: upper bound the minimax regret for arbitrary expert classes.
Difficulty: logarithmic loss is neither bounded nor Lipschitz.

Sequential Covering

Cesa-Bianchi & Lugosi (1999) use a uniform covering of $\mathcal{F}$. This is too coarse for many expert classes.

Similarly to Rakhlin & Sridharan (2015) and Foster et al. (2018), we rely on *sequential covering*, introduced by Rakhlin & Sridharan (2014).

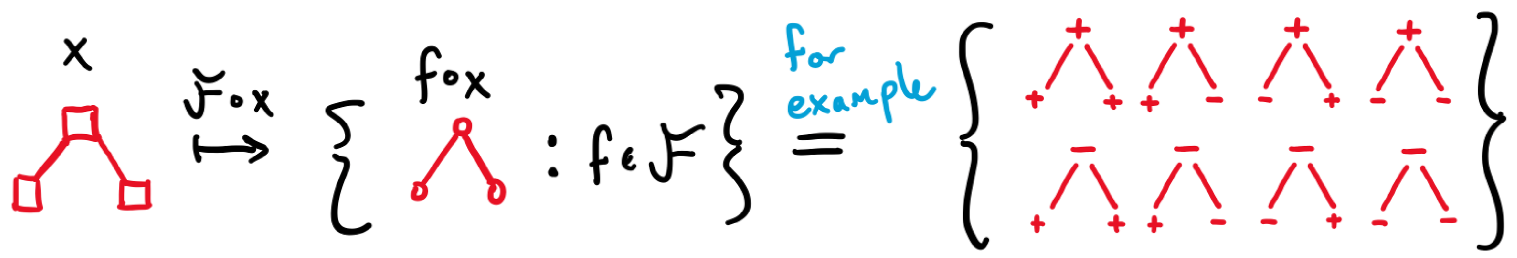

Fig: Composition of context tree with experts illustrated for binary experts.

An exact sequential cover of the binary classification example requires only 4 trees rather than the 8 needed for a uniform cover, since **a new covering element can be chosen for each path** rather than only for each tree of $\mathcal{F} \circ \mathbf{x}$.

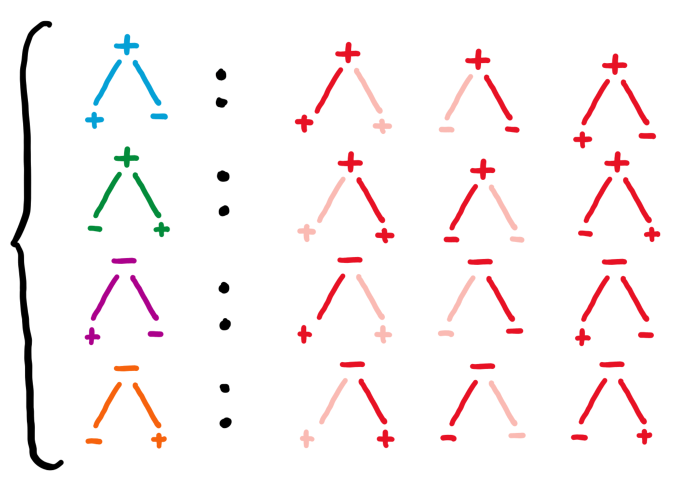

We denote the sequential γ-covering number by $\mathcal{N}_\infty(\mathcal{F} \circ \mathbf{x}, \gamma)$.

Improved Upper Bound

For any context space $\mathcal{X}$ and class of experts $\mathcal{F} \subseteq [0,1]^{\mathcal{X}}$:

$$R_n^{\log}(\mathcal{F}) \le \sup_{\mathbf{x}} \inf_{\gamma > 0} \left\{ 4n\gamma + c \log\left(\mathcal{N}_\infty(\mathcal{F} \circ \mathbf{x}, \gamma)\right) \right\}, \qquad (1)$$

where $c = \frac{2 - \log(2)}{\log(3) - \log(2)}$.

Ask me why this bound does not use chaining.

Sequential Covering Number Examples

- **Time-Invariant:** $\mathcal{F} = \{ f(x) = q \, \forall x \in \mathcal{X} \mid q \in [0,1] \}$.
  $$\sup_{\mathbf{x}} \log\left(\mathcal{N}_\infty(\mathcal{F} \circ \mathbf{x}, \gamma)\right) \le \log(1/\gamma).$$
- **1-Lipschitz:** $\mathcal{F} = \{ f : \mathbb{R} \to [0,1] \mid |f'(x)| \le 1 \}$.
  $$\sup_{\mathbf{x}} \log\left(\mathcal{N}_\infty(\mathcal{F} \circ \mathbf{x}, \gamma)\right) = 1/\gamma.$$
- **Linear Predictors:** $\mathcal{F} = \{ f(x) = \frac{1}{2}[1 + \langle w, x \rangle] \, \forall \|x\| \le 1 \mid \|w\| \le 1 \}$.
  $$\sup_{\mathbf{x}} \log\left(\mathcal{N}_\infty(\mathcal{F} \circ \mathbf{x}, \gamma)\right) = 1/\gamma^2.$$

Comparison to Previous SOTA

We compare our upper bound from (1), denoted $U_n^{\text{new}}(\mathcal{F})$, to the previous best upper bound from Foster et al. (2018), denoted $U_n^{\text{old}}(\mathcal{F})$. For any context space $\mathcal{X}$ and class of experts $\mathcal{F} \subseteq [0,1]^{\mathcal{X}}$:

1. If $\sup_{\mathbf{x}} \log\left(\mathcal{N}_\infty(\mathcal{F} \circ \mathbf{x}, \gamma)\right) \le \mathcal{O}(\text{polylog}(1/\gamma))$,
   $$\frac{U_n^{\text{new}}(\mathcal{F})}{U_n^{\text{old}}(\mathcal{F})} \le \mathcal{O}\left(\text{polylog}(n)\right).$$
2. If $\sup_{\mathbf{x}} \log\left(\mathcal{N}_\infty(\mathcal{F} \circ \mathbf{x}, \gamma)\right) \asymp 1/\gamma^p$ for $\boldsymbol{p \le 1}$,
   $$\frac{U_n^{\text{new}}(\mathcal{F})}{U_n^{\text{old}}(\mathcal{F})} \le \mathcal{O}\left(\frac{1}{\text{polylog}(n)}\right).$$
3. If $\sup_{\mathbf{x}} \log\left(\mathcal{N}_\infty(\mathcal{F} \circ \mathbf{x}, \gamma)\right) \asymp 1/\gamma^p$ for $\boldsymbol{p > 1}$,
   $$\frac{U_n^{\text{new}}(\mathcal{F})}{U_n^{\text{old}}(\mathcal{F})} \le \mathcal{O}\left(\frac{1}{n^{\frac{p-1}{2p(p+1)}}\text{polylog}(n)}\right).$$

Self-Concordance

Our proof technique exploits the *self-concordance* of logarithms.
A function $F : \mathbb{R} \to \mathbb{R}$ is self-concordant if for all $x \in \mathbb{R}$,

$$|F'''(x)| \le 2F''(x)^{3/2}.$$

Ask me about this, and how it leads to a truncation-free argument.

Lower Bound

If $p > 0$, there exists an $\mathcal{F}$ with $\sup_{\mathbf{x}} \log\left(\mathcal{N}_\infty(\mathcal{F} \circ \mathbf{x}, \gamma)\right) \asymp \gamma^{-p}$ and

$$R_n^{\log}(\mathcal{F}) \ge \Omega\left(n^{\frac{p}{p+2}}\right).$$

NSERC CRSNG